Scalable algorithms for Markov process parameter inference

Darren Wilkinson, Newcastle University

Inferring the parameters of continuous-time Markov process models using partial discrete-time observations is an important practical problem in many fields of scientific research. Such models are very often "intractable", in the sense that the transition kernel of the process cannot be described in closed form, and is difficult to approximate well. Nevertheless, it is often possible to forward simulate realisations of trajectories of the process using stochastic simulation. There have been a number of recent developments in the literature relevant to the parameter estimation problem, involving a mixture of approximate, sequential and Markov chain Monte Carlo methods. This talk will compare some of the different "likelihood free" algorithms that have been proposed, including sequential ABC and particle marginal Metropolis Hastings, paying particular attention to how well they scale with model complexity, and their parallel implementation in modern functional programming languages.